

EXHIBIT 3

NETWORKWORLD Reprint

THE CONNECTED ENTERPRISE = FEBRUARY 22, 2013

How Arista Networks got out in front of the SDN craze

Arista CEO Jaysree Ullal says 'cloud networking leader' complements Cisco

BY JOHN GALLANT, NETWORKWORLD

Today, the buzz in networking is all around software-defined networks — and nothing could make Arista Networks CEO Jaysree Ullal happier. Ullal spent 15 years at Cisco, where she ran the network giant's core switching and data center businesses, before joining Arista, which was founded by Sun Microsystems co-founder and Chief System Architect Andy Bechtolsheim and David Cheriton, a Stanford University professor of computer science and electrical engineering (and fellow Cisco alumnus). Ullal says Arista's data center switches were born to support SDN and provide both the power and flexibility required for today's highly virtualized corporate and cloud data centers. In this installment of the IDG Enterprise CEO Interview Series, Ullal spoke with Chief Content Officer John Gallant about the reality and hype around SDN, and why the data center requires a different network than your father's general-purpose Cisco net. She also explored how her work at Cisco shaped Arista's strategy, and shared insights on how Arista's partnerships with VMware and Cloudera are making it easier to move to cloud and embrace big data, respectively.

There are a lot of networking alternatives out there. Why should someone buy from Arista?

Arista saw three disruptions in the market: a hardware disruption; a software disruption; and a customer buying disruption, which in my mind is the most important thing: You can invent all you want on the technology side, but you have to see the customers changing their market position.

The hardware technology disruption was that in the 1990s, the only way to build any kind of high-speed networking was through your own in-house ASICs [application-specific integrated circuits] and specialty chips. That's not true anymore. We have from three to five vendors available, whether it's Intel, Broadcom or others, supplying us much of the silicon. They are sometimes an order of magnitude better in power, footprint, density, latency, and performance and scale.

Arista was able to take advantage of that disruption in hardware.

The second is software. We were very inspired by Cisco's software focus on the enterprise side, Juniper's on the service provider side, and we saw that we could build a purpose-built, modern operating system only for the data center and the cloud. We didn't try to do it for general-purpose networking. We really focused on our mission, which is high-performance applications for the data center and cloud. It's called Extensible Operating System (EOS) and there is no networking operating system that is as modern, self-healing and resilient, and [designed for the cloud].

And the third, speaking of that, is the cloud itself. The enterprise market is shifting. Every CIO is being demanded a strategy on what they are doing with the cloud in terms of applications and infrastructure. Whether it's a private cloud, a public cloud or a hybrid cloud, these are becoming an important piece of the strategy. As Amazon innovated on the application side, you can think of Arista as really providing that market disruption on the networking side.


Explain the cloud angle in a little more depth. What were you setting out to do to support or enable cloud?

More and more people are outsourcing to modern applications — whether it's Salesforce.com or Amazon itself. [They're supporting] high-performance computing, or high-frequency trading or, increasingly now, big data and network virtualization. The network infrastructure needs to adapt. It cannot be so monolithic. It cannot be one physical port equals one VLAN equals one network switch. It really needs to be much more massive in scale. A typical enterprise network is a 10,000-node, three-tier network, and we were able to build a much flatter, fatter topology at Layer 2 and 3, using what we call the leaf-spine architecture that can scale to 50,000 to 100,000 nodes. That was our first premise.

The second [thing we focused on] was application delays. Don't build a network

Our top five differentiators are all tied to our software."

— Jaysree Ullal, CEO,
Arista Networks



as a cost center, but really build it as a profit center by addressing the applications themselves. We early on entered the high-frequency trading market to understand their trading algorithms, map it to the latency requirements. That became an instance of a high-performance financial cloud where they started building the network for that application separate from the enterprise network.

In Silicon Valley, a large number of Web 2.0 providers, whether they're search engines or social networking, the kind of scale they build is just unbelievable. It's 100,000 nodes, and increasingly, one machine, one physical server, is not one node. That's 20 virtual machines, which means you could be enabling 100,000 physical nodes but you are really enabling 1 million virtual nodes. There's huge virtual machine sprawl and physical sprawl. The CPU at one point wasn't being fully utilized. But now, with the new multi core CPUs, the pressure is back on the network. That's why whether it's a private or a public cloud, the Web 2.0 companies are moving massively to high-density 10G, 40G and 100G [networks] that are requiring a new type of architecture and new software as well.

What are the things that make you different than a general-purpose networking company like Cisco?

At the highest level I would say our software, our EOS. It's open, it's built out of straight Linux. But then we added what we call multi-processing, state-oriented software that allows you to do the kind of things that you could only do in mainframes and servers. It's funny how hardware changes every 18 months in networking, but software doesn't change for decades and has remained monolithic

for so long. Our top five differentiators are all tied to our software.

The first is that we build, without using any proprietary components, active/active networks that can scale to 50,000 and 100,000 nodes. Other companies try to do that with proprietary technologies. You may be aware of Juniper's QFabric or Cisco's FabricPath and OTV [Overlay Transport Virtualization]. We are able to do it in a standards based fashion, and every one of our networks interoperates with Cisco routers, Juniper switches, NetScreen firewalls, you name it.

The second is, because of the software, we were able to bring to the data center and cloud what we call self-healing resilience. Usually, redundancy and resilience means buy two of everything and connect them in case one fails. It's great for the vendor to get two of everything. But we were able to do it right in our software. Today, you look at software agents and how they interact. If you have a memory leak in software today, and the agents talk to each other in a traditional network operating system, they do so with something called IPC, inter-process communication. But think of the cloud where you have, like we described, 100,000 of these, the multiplier effect of failure is huge with this inter-process communication. Arista chose a publish/subscribe model using a built-in SYSDB database, where the state of every software agent is stored. Because that's not human-generated, it's the most resilient piece of code. Let's say you have a failure. We automatically track the failure and contain it. Then we repair it. We actually spin up a new agent. Today's enterprise agent manager has no maintenance windows. So they don't have to know.

The third [differentiator] is that we are open and programmable. You hear a lot of talk about SDN these days, and one has to separate the hype from the reality. The essence of SDN to me is, first of all, build open interfaces and allow your customers to write to their applications through our APIs at the northbound level, and at the southbound level our devices must be programmable. We didn't call it SDN back when we developed this, we called it EOS. The extensible in EOS is [in reference to the operating system being] very programmable. Every aspect of our software, whether it's at the hardware plane, at the device plane or the software plane, can be programmed. That's a huge advantage. We find ourselves in a fortunate position that as the SDN market is evolving, our network is already open and programmable and SDN-ready.

The fourth one is big data analysis. Data analysis and traffic visibility is becoming a real weakness, because, as you know, we can all talk about improving price, performance and CAPEX, but the biggest cost center in networks is OPEX. There are three ways to solve OPEX issues: Stop buying gear, outsource your gear or make your technology do better work. We believe technology to solve the problem is far better than outsourcing or throwing people at the

problem. We call this "from A to Z analysis." We can do automation, we can do zero touch provisioning, we can do a suite of functions here because data is coming at such amazing speeds, structured and unstructured, how do you sort out what's relevant and how do you monitor, how do you tap, how do you do real-time captures at 10 gigabits and terabits when the data is moving so fast? We're not just building enterprise features. Cisco's done that really well for the last two decades, that's their market. But yet if you look at the way servers are sold today, only half of them are going into an enterprise application. The other half, which are high-performance computing and Web, are going into the cloud applications. They don't require traditional enterprise features. Just like mainframes moved to client server, enterprises are moving to more HPC and Web, and those features are much more about reducing OPEX and improving the orchestration and traffic visibility and data analysis.

The fifth and final differentiator is network visualization. What VMware did to servers with server virtualization, we believe jointly working with VMware we can do with network virtualization. VM sprawl has created network sprawl. Arista and VMware, together with a number of other vendors, Broadcom, Cisco, etc., defined to me what is one of the most breakthrough specifications in our industry — VXLAN, virtual extended LAN. The VLAN, as a unit, is something we all grew up with and invented back in the '90s. It's been with us 25 years, way too long. VLAN boundaries have plagued the deployment of virtualization because you're limited to 6,000 VLANs or 16,000 VLANs, and you've got many more virtual machines. So therefore, you've had a vi-admin manage one, the virtual network, and the command line interface or Cisco admin manage the physical network. These two worlds need to come together. Arista, working particularly closely with VMware, has been able to bridge that gap between network physical and network virtual, using VXLAN. VXLAN all of a sudden opens up the boundary from 16,000 to 16 million possible entries. So we're very excited with the technology we demonstrated at [the VMworld conference].

Is it deployed now in the market?

Very early. We are one of the first to come out with it. We showed it August 2012, and we showed interoperability with VMware, EMC and F5. We shipped a product based on it, the Arista 7150, in November.

Say I'm a big Cisco installation today. When would I talk to Arista? What's the need that opens the door?

It could be project-based or it could be a strategy. When it's project-based, it's usually that you're deploying high-frequency trading or you need a high-performance compute solution, usually InfiniBand and Ethernet get reviewed. Sometimes InfiniBand gets chosen because the supercomput-

er guys really like it and other times it's high-density 10GB. Another application is big data. Storage is no longer just a fibre-channel SAN — you will start needing 10GB storage for iSCSI or more and more Hadoop clusters with direct-attached storage. That becomes another very interesting Arista project. Virtualization, the VM sprawl. Another one we're starting to see more of is huge media rendering, and video applications that are pushing the envelope of bandwidth. Where the application intersects the network is the common theme through all the projects.

On the other hand, Arista has to walk before it runs. We've been growing at the rate of one new customer a day since we started shipping. We now have 1,700 customers. Deployments usually start small, then they get really fascinated and intrigued and appreciative of EOS, and all of its operational advantages, how open it is, how easy it is to use. The training is very easy and a Cisco CCIE expert would be able to use Arista right away, because we have similar command-line interfaces and operational look and feel. Where we don't have to invent, we don't. Where we had to invent for these specific use cases we do, so most often it's a use case or a project. Sometimes it's a data center build-up. After they use us in one project, they'll say they want to consolidate data centers. I would say 10% to 20% of them are now standardizing on Arista as their data center strategy.

You mentioned about 1,700 customers. Give me a sense of your business progress to date.

We're not supposed to [talk revenue] but the company is very young, it's only 5 years old. We've gone from 30 employees when we started to more than 500. I guess the biggest thing I'd leave you with is that in the beginning we were a market leader for financials and high-frequency trading which, as you know, is a tough customer. We've always had to go into mission-critical [environments] and we didn't have it easy ever. It wasn't like we were in a little lab somewhere. We believe we are today 70% to 80% market leaders in high-frequency trading. In 2008, 2009 and even part of 2010, that was 70% of our business. Today it's diversified nicely into three areas. The first is financials. The second is what I call Web 2.0, and the massive scale of their deployments, the cloud scale, really. The third is cloud and service providers. Every service provider is looking to be a cloud vendor. In all of these three cases we are being looked at as the innovative alternative to traditional legacy players.

You were at Cisco for a long time.

Yeah, 15 years. I intended to be there two years. But I was there 15 years, two years at a time.

So how did your experience at Cisco shape this?

I had a big hand in shaping Cisco's enterprise switching strategy, and it helped me

appreciate what to do and what not to do. At some level I don't feel we compete with Cisco because we're not taking on the traditional enterprise market. But at other levels I feel like I learned a lot about what not to do and how to do [things] better, by being more application-focused and really taking advantage of the market disruption from enterprise to cloud, and then to big data and network virtualization. Cisco, in my view, will always be the enterprise market leader. Arista is inspired and aspires to be the cloud networking leader and really be a complement to Cisco.

Other networking companies have tried to address the needs you have described. Let's take a case in point: Juniper seemed to be targeting the same kinds of problems, so why did the company struggle with its QFabric?

Without making it specific to one vendor, I would say three or four years ago, you rightly pointed out that the market was very crowded. There was Juniper, HP looked like they were coming on strong, there was Force10 that then got acquired by Dell, and then there was Brocade that had acquired Foundry. They have taught me that focus is important. We stayed relentlessly focused on building a standards based open architecture for data centers. I think many of these vendors, they get distracted by growth. It's a difficult call, in the public market especially. It's easy to do when you're private and you don't have to explain things to people. But do you go broad and wide or do you go a mile deep? My belief is I am participating in one of the largest total available markets in the history of networking. The 10G market is going to go from a couple of billion to \$15 billion in 2016. I should not be distracting myself with other markets. I should relentlessly focus on doing my one market very well. I would say that is the failing of many competitors who don't have that kind of focus.

How has the strategy changed since launch? What have you learned in that time period?

We built a point product then and our strategy has changed to a portfolio. We have gone from a top-of-rack (48 ports) to 384 ports in a chassis. Secondly, our software has evolved, not just in millions of lines of code, but in terms of the agility and innovative features. We've been putting out releases practically every quarter since then, to the point that some of our customers say: "Hey, slow down a little! We're unable to absorb it as quickly as you put it out." One customer said to us: "We have all these data centers and we're racking up our servers, we're connecting it to a VLAN and we're enabling DHCP, this whole process is taking us two hours." With our zero-touch provisioning we were able to cut that down to 20 minutes. You think of this in one data center and how

you replicate it, it's a huge multiplier. The third area I'd say we have really evolved is partnerships. We've gotten closer to the big data companies like Cloudera and the virtualization companies like VMware. Because we are best-of-breed, we are in some ways less of a threat and more of a partner to many. The security companies, like Palo Alto Networks, the application delivery companies like F5, view us as a friendly face.

What do you do with Cloudera?

We have actually installed joint networks for big data together. The biggest issue with using these kinds of direct-attached Hadoop systems is that you have to have a network with fast failover characteristics, the right buffering characteristics, and you actually almost have to have a Hadoop tracer-like function between the storage and the network. Because, remember, all of a sudden you wiped out the concept of a storage-area network, but you still need to have the resilience of a storage-area network. We're in several joint customers together, particularly in mission-critical financials.

I want to go back to EOS for a bit. There is sort of this myth of the unified operating system in networking, but competitors and customers are all running multiple versions and flavors of these OSes. What makes this software different? Why would somebody believe that you have a single unified network operating system where people have struggled with that?

We've proven it. Even though we've been here five years and we've done all the software releases, we still have one single binary image. Nobody has to get a Ph.D. on our software releases because we've kept one single unified code base. Secondly, when people play with the software, they realize that: "Oh, this can do Python scripting, I can write agents to it." So the power that we've given them to enable things that even we don't enable is like no other. Now, I could tell you some people love it. The engineering community loves it. The CIO community is afraid of it. So we also have a lock-down mode where we can have all the security and people can't go in and simply start writing.

At the same time, I would say we are a little bit of a Sybil with our EOS. One side of us looks just like a Red Hat or Linux, and the server guys have fun with it. They can do TCP and dump and bash. It's a Linux kernel, right? Then the other side, for a Cisco administrator, looks just like a Cisco network. When we work with VMware, we talk straight into vCenter and vSphere. So I think the flexibility that we have offered in being open, at the same time not destroying the paradigm between the virtual admin and the sys admin for the server, and the network admin has been unique.

Let's go back to SDN. What's real and what's false about SDN?

If you ask 10 people what SDN means they'll give you 10 different answers. But if I had to describe it in one or two words, I would say: open and programmable. There's been so much vendor lock-up in networking, with the huge operational cost of being locked in with one vendor. That mainframe model in networking is what SDN is challenging. Why is SDN fever and hype so high? People are sick and tired of the vendor lock in and proprietariness and they are looking for a movement. Now, how and when will that movement happen? Like anything, you have to be pragmatic about that. When the hype is so high the market isn't that big, right? But in my view, SDN has a tremendous opportunity to succeed if we're pragmatic about the use cases. Is SDN OpenFlow? No, it's not just OpenFlow. Is it OpenStack? No, it's not just OpenStack. Is it VXLAN or network virtualization? Understanding what you can actually do with SDN is the key here.

Let's talk about that programmability aspect. When I talk to people in the market, there are some who are really excited about that piece, but others who think there are only a limited number of things you would ever want to program the network to do. What do you envision people programming the network to do?

I think they're both right for different reasons. In a traditional network, no one is looking to toss their IP out any time soon, so whatever you do you've got to make sure your IP network is up and running. But say you're going to do a green field [installation]. This is how the OpenFlow SDN movement got started at Stanford. [Professor] Nick McKeown was doing this project called Clean Slate, where he was telling his students: "Imagine a world with no IP. How would we define it?" That's all great for vision and strategy, but we've got hundred of thousands of customers operating networks, so you have to understand that no matter how inflexible and how much of a headache your current architecture may be, TCP/IP does work. Then you look at the model that SDN is coming up with, which I call the controller, overlay controller model. That's a controller for OpenStack, OpenFlow, network virtualization, each of them is a use case — it's a specific case where you need programmability. I agree that you shouldn't go and mess with the big IP network. If it's not broken, don't fix it.

Understand what are the use cases you're trying to augment with programmability. I can think of three or four use cases. One of the ones that we found, especially with OpenFlow, but also in IP networking, is data tap aggregation. When you're running at the 10G speeds we are, everybody is looking for traffic visibility and understanding what's going on in the network. You can build an out-of-band

controller with OpenFlow, whether it's Big Switch Networks or open source controllers from Floodlight or NEC, and have an OpenFlow agent on our switches, and have a very simple SDN network that's highly programmable, and still works in hybrid mode with your IP network. That could be one case. There's the Nicera case, which is also, in my view, a programmable use case, but it's for strictly network virtualization. You keep the IP running, but you need a network virtualization platform that can program your virtual switches, whether they're VMware switches or open virtual switches from an OpenStack environment. Today, this is literally like two islands; there's the physical switch and the virtual switch. With Arista working closely with VMware and Nicera, we can transcend the virtual-to-physical islands, where every VXLAN vSwitch port automatically maps to a hardware port. Now you've got network virtualization not as two separate failure domains, but transcending virtual-to-physical, potentially even to a cloud, architecture down the road. This is in a VMware environment, but there's no reason you couldn't do it in an OpenStack environment with quantum plug-ins as well. So that's another use case.

Do you believe SDN fundamentally changes the competitive landscape?

Yes, but I believe the controller vendors in isolation won't succeed. The networking vendors, if they get defensive with just IP won't succeed. You have to have the Arista view, which is have two personalities. Work in a controller mode, but work also with your IP network. That's the mistake I think SDN is making, in that they're thinking of it as only green field. In reality there's a whole world of IP there that you need to work with while you're trying to develop the use cases.

Server virtualization was essentially owned by VMware and the software vendors, not the server vendors. Who ends up owning SDN?

[The server vendors] got defensive. If the networking vendors do not embrace it because it requires a new software paradigm that they haven't built, then they stand to lose, because over time there will be more and more use cases that disrupt them. Customers like that. I was talking to a customer today who said: "The most liberating feeling for me now is I can be multi-vendor, and it hasn't cost me that much money.

I don't know if there are that many SDN use cases, but I love it." Arista feels very fortunate that we got an eight-year head start. We spent four years building the software, four years commercializing it, and now we're sitting in the middle of an SDN momentum and a switch momentum that puts us on the natural cusp, if you will, where we can be working with the old and still developing the new.

Does Arista have its own controller?

No.

Is that something you're going to have?

Our control plane and data plane will be open, and we will work with all the controller vendors that you deploy. HP Opware, vSphere VMware, Nicera, Big Switch, Floodlight, NEC, IBM Tivoli, EMC Smarts. These are all to us controllers. We don't view ourselves as the management expert. Now, if you want to go develop specific use cases for a controller on our switch, we can do that too. And people do that. Like for example, we worked with Splunk very closely to develop a configuration management tool with them.

If you're a CIO or a senior IT executive, and you are on the sidelines looking at SDN and wondering what it's all about, what's the 30-second answer on why they should care about it and start thinking about it now?

Because it unshackles their traditional networking decisions from their application and helps them develop new use cases that they couldn't do before, or only had to do with traditional vendors. The key for them is to understand what problem they are trying to solve. Solve that out-of-band and work with the existing [network].

Using merchant silicon versus custom ASICs, is there a limit to where that can take you?

I think custom ASICs have fallen behind by several years, and unless a vendor can build and compete against merchant silicon, there's no point in doing custom ASICs. It's something that the established vendors have to come to terms with, just like the server community did. There was a time when SPARC was sparkling for Sun, and then everybody had to embrace the x86 Intel world. A similar phenom-

enon and revolution is happening in networking.

But does it limit the speed at which you can develop your product?

No, actually, it's better. We were most scared when we were only beholden to one vendor. The way our software architecture works, 90% of our engineering is all software. Then there's a driver layer. The driver layer is where we customize to the silicon. Today we operate with Intel drivers, with Falcon drivers, with Broadcom drivers, we used to have a company called Dune that got acquired by Broadcom, you know, Marvel. The more drivers, the more we have the roadmap and the R&D of all the silicon vendors, because remember there are always risks. The biggest risk with the silicon vendors is if they miss a cycle. They don't miss it by one or two weeks. They often miss by one or two years, whether they work for you inside the company or outside. We like the fact that we have drivers and have developed drivers over time for all these chip vendors.

What's ahead for Arista in 2013?

Expect to see more capacity, more software, more of everything, and the replacement of our older products. Second, expect us to continue to be application focused, because that's where we think we get appreciated. If it's a straight, traditional network, where nobody gets fired for buying IBM or Cisco, that's not where Arista shines. Arista will continue to provide thought leadership, like we did with high-frequency trading, and like we do now with big data and network virtualization. We're very excited when we think of the number of things we have to do in the software and application world, and prioritizing them, there's so much ahead of us.

Last question: the classic elevator situation. You've got 15 floors with a CIO. You tell them what about Arista?

We're the best-of-breed data center and cloud networking alternative in the market. We're not a cost center, we are an enabler of applications. There are people in high-frequency trading who have told us: "We make \$2 million a day on your boxes." In a nutshell, that's what I would say. We are the thought leader in cloud networking, we enable the applications and we save you CAPEX and OPEX.

ARISTA

R&D and Headquarters

5470 Great America Parkway • Santa Clara, CA 95054 • Tel: 408-547-5500

Email: info@aristanetworks.com

US & North America Sales: us-sales@aristanetworks.com

Latin America Sales: latam-sales@aristanetworks.com

Europe, Middle East & Africa Sales: emea-sales@aristanetworks.com

Asia-Pacific Sales: apac-sales@aristanetworks.com

Japan Sales: japan-sales@aristanetworks.com